Large-scale Stabilized FE Computational Analysis of Nonlinear Steady State Transport/Reaction Systems[1]

J. N. Shadid, A. G. Salinger, R. P. Pawlowski, P. T. Lin, G. L. Hennigan, R. S. Tuminaro, R. B. Lehoucq

Sandia National Laboratories, PO Box 5800, MS 0316, Albuquerque, NM 87185-0316

Abstract

The solution of the governing steady transport equations for momentum, heat and mass transfer in fluids undergoing non-equilibrium chemical reactions can be extremely challenging. The difficulties arise from both the the complexity of the nonlinear solution behavior as well as the nonlinear, coupled, non-symmetric nature of the system of algebraic equations that results from spatial discretization of the PDEs. In this paper, we briefly review progress on developing a stabilized finite element (FE) capability for numerical solution of these challenging problems. The discussion considers the stabilized FE formulation for the low Mach number Navier-Stokes equations with heat and mass transport with non-equilibrium chemical reactions, and the solution methods necessary for detailed analysis of these complex systems. The solution algorithms include robust nonlinear and linear solution schemes, parameter continuation methods, and linear stability analysis techniques. Our discussion considers computational efficiency, scalability, and some implementation issues of the solution methods. Computational results are presented for a CFD benchmark problem as well as for a number of large-scale, 2D and 3D, engineering transport/reaction applications.

[1] This work was partially funded by the Department of Energy Office of Science MICS program at Sandia National Laboratory. Sandia is a multiprogram laboratory operated by Sandia Corporation, a Lockheed Martin Company, for the United States Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.

1 Introduction

Physical problems in a number of scientific and engineering fields can be described by a system of coupled nonlinear PDEs. One of the main interests in these fields is the problem of determining steady and/or time dependent states evolving with changes of characteristic parameters. These parameters arise in various ways. The parameters can be coefficients of the PDEs, characteristic values from the boundary conditions, or characteristic dimensions from the domain of the solution. Our discussion focuses on applications involving steady state solution of transport/reaction systems. However, the methods, algorithms, and parallel implementations presented have broad applicability to many scientific fields.

The numerical challenge for transport/reaction system simulation is the solution of the partial differential equations (PDEs) describing momentum, heat, and multi-component mass transfer with chemical reaction source terms. These governing PDEs are outlined in residual form in Table 1, and are valid for low Mach number flows operating at a thermodynamic system pressure of $\tilde{P}$. In these equations the transport properties are the mixture density, ρ, mixture viscosity, μ, mixture specific heat, $\hat{C}_P$, the mixture thermal conductivity, λ, the k^{th} species mass diffusivity, D_k, and thermal diffusivity, D_k^T. The source terms are defined by the molecular weight, W_k, the specific enthalpy, $\hat{h}_k$, and the molar production per unit volume, $\dot{\omega}_k$ for the k^{th} species. This system of PDEs is non-self adjoint, strongly coupled, highly nonlinear, and characterized by physical phenomena that span a large range of length and time scales. The high degree of coupling and nonlinearity in this system is generated from the convection terms, the chemical reaction source terms, and the strong dependence of the transport properties and chemical kinetics on the thermodynamic state, $(T, \tilde{P}, Y_k)$.

Our discretization of the governing transport/reaction equations employs stabilized finite element (FE) methods. This formulation is based on the developments of Hughes et. al.(see e.g. [35,5,34,32,39,31,37,36,67]) in a seminal set of papers that developed this very successful and powerful discretization methodology. The stabilized formulation circumvents the Ladyzhenskaya-Babuska-Brezzi condition (see e.g. [4,25]) for compatible discretization for mixed finite element formulations of the saddle point problem arising from discretization of the incompressible Navier-Stokes equations. The stabilized FE formulation allows for equal order interpolation of the incompressible Navier-Stokes equations and eliminates spurious pressure modes. Our formulation implements a simplified form of a consistently stabilized FE method [1]. In the current context of transport/reaction systems, the use of equal order interpolation simplifies the data structures of a parallel unstructured FE code and the linear algebra interface for iterative solution methods.

Momentum	$\mathbf{R}_m = \rho(\mathbf{u}\cdot\nabla\mathbf{u}) - \nabla\cdot\mathbf{T} - \rho\mathbf{g}$
Total Mass	$R_P = \nabla\cdot(\rho\mathbf{u})$
Thermal Energy	$R_T = \rho\hat{C}_p\left[\mathbf{u}\cdot\nabla T\right] + \nabla\cdot\mathbf{q} + \sum_{k=1}^{N_s}\hat{h}_k W_k\dot{\omega}_k$
Mass Fraction for Species k	$R_{Y_k} = \rho\left[\mathbf{u}\cdot\nabla Y_k\right] + \nabla\cdot\mathbf{j}_k - W_k\dot{\omega}_k, \quad k = 1,2,...,N_s-1$ $R_{Y_{N_s}} = \sum_{k=1}^{N_s} Y_k - 1$
Newtonian Stress Tensor	$\mathbf{T} = -P\mathbf{I} + \mathbf{\Psi} = -P\mathbf{I} - \frac{2}{3}\mu(\nabla\cdot\mathbf{u})\mathbf{I} + \mu[\nabla\mathbf{u} + \nabla\mathbf{u}^T]$
Heat Flux Vector	$\mathbf{q} = -\lambda\nabla T + \sum_{k=1}^{N_s}\hat{h}_k\mathbf{j}_k$
Mass Species Flux Vector	$\mathbf{j}_k = \rho D_k\nabla Y_k - D_k^T\frac{\nabla T}{T}$

Table 1
Residual form of governing transport/reaction PDEs and simplified constitutive equations. The primitive variables are the velocity vector $\mathbf{u}$, the temperature T, the hydrodynamic pressure P, and the N_s species mass fractions Y_k. The physical and transport properties are a function of the local T and Y_k and a global system pressure $\tilde{P}$.

Additionally, a stabilized FE strategy is also used to control instability in the Galerkin FE formulation for highly convected flows. The methodology that we employ is based on a variation of the stabilized FE formulations of Hughes et. al. [34,39] , Shakib [65,66], and Tezduyar [70]. The stabilized FE method allows solution of systems where the cell Reynolds number, Re_c, and thermal energy and mass transport Peclet numbers, Pe_c, are larger than one by decreasing numerical oscillations due to convection. In addition, this stabilization improves the conditioning of, and therefore the iterative solution of, the Jacobian matrices in the linear subproblems generated by Newton's method.

In general, solutions of the transport/reaction PDEs described in Table 1 can produce complex nonlinear solution behavior. Through the variation of one or more parameters, new multiple steady states, time periodic, quasi-

periodic, or chaotic (aperiodic) solutions appear (or bifurcate) at critical values of the parameters [45,24,12]. Depending on the path followed through the parameter space, these multiple solutions can appear, become more complex and disappear with the change of the parameters. Not only is the number and location of these solutions in the parameter space of interest, but also the stability or long time evolution of such systems. The existence of an instability (growth of disturbances) in a physical or engineered system is of great concern, because it often results in defects or even breakdown of processes. It is through an understanding of the physical mechanisms that stabilize and destabilize such systems that active control and process optimization of these systems is possible.

The ability to analyze computationally the steady state behavior of many complex engineering systems centers on the ability to quickly and reliably obtain solutions to nonlinear algebraic systems with millions of unknowns solved on large-scale parallel machines. To do this we employ parallel Newton-Krylov solution strategies [6,20,43,54] that are robust and scale well on distributed memory architectures. These solution methods are critical kernels used by continuation methods for tracking the solution branches with respect to key system parameters. Specialized methods such as Euler-Newton and pseudoarclength continuation techniques can be used to follow regular solution paths, unstable solutions, and critical bifurcation structures in nonlinear solutions. Many steady state solutions lose their stability to symmetry breaking (pitchfork bifurcations), turning points, or periodic oscillations (Hopf bifurcations). The stability of the solution branches to infinitesimally small disturbances is determined by the use of linear stability theory [50,11] . A linearizion of the nonlinear system produces a generalized eigenvalue problem from which the stability of the nonlinear system is determined. The solution of the generalized eigenvalue problem is difficult when dealing with large non-symmetric systems of equations [55]. For this reason, special techniques must be used to efficiently determine the eigenvalues with algebraically largest real parts that govern the stability of the solution branch. One such technique is to recast the eigenvalue problem using the Cayley transformation, and then employ an Arnoldi method to approximate the eigenvalues and eigenvectors [49].

The remainder of this paper is organized as follows. The stabilized FE formulation of the governing transport/reaction equations is presented in Section 2. In Section 3 a brief overview of the solution of large sparse nonlinear algebraic systems with parallel Newton-Krylov techniques is presented. Parameter continuation and bifurcation tracking methods are presented in Section 4. The linear stability problem along with the solution of the non-symmetric generalized eigenvalue problem is discussed in Section 5. In Section 6 we present some representative performance, scaling and simulation results of these solution methods for some illustrative transport/reaction systems. Finally in Section 7 we close with a few conclusions.

2 Stabilized FE Formulation of Governing Transport/Reaction Equations

2.1 Brief Overview of Stabilized Equations

Table 1 presents the conservation equations for momentum, total mass, thermal energy and mass species transport with non-equilibrium chemical reactions for low Mach number flows. Table 1 includes a specific description of constitutive equations for a Newtonian variable density fluid mixture with a heat flux vector with thermal conduction and diffusional mass transfer contributions. The diffusional mass species flux vector has transport due to both concentration gradients as well as thermal diffusion (Soret) effects. The reacting flow examples that are presented in this study are based on non-equilibrium statistical mechanical theory of multicomponent, dilute polyatomic gases. In this case, necessary transport properties, diffusion coefficients, kinetic rate constants and diffusion velocities are obtained from the CHEMKIN [41] subroutine library. This library provides a rigorous treatment of dilute-gas multicomponent transport, including the effects of thermal diffusion. Chemical reactions occurring in the gas phase and on surfaces are also obtained through CHEMKIN. In general, the necessary constitutive relations for the stress tensor, $\mathbf{T}$, heat flux vector, $\mathbf{q}$, and species mass fluxes, $\mathbf{j}_k$, can be quite complex. A detailed description of more general transport/reaction equations and chemical kinetics formulations can be found in [60].

The continuous PDE problem, defined by the transport/reaction equations in Table 1, is approximated by a stabilized FE formulation ([35,5,34,39,31,37,36,67]). This formulation allows for stable equal order velocity-pressure interpolation and provides for convection stabilization as described above. The resulting stabilized FE equations are shown in Table 2. The stabilization parameters (the τ's) are based on the formulations of Hughes and Mallet [38], Shakib [65,66], Hughes [33], and Tezduyar [70].

The definition of the stabilization parameters are provided in Table 3 for momentum, energy, and mass species. The multidimensional effect of convection is incorporated into the stability parameters by the use of the contravarient metric tensor, $\mathbf{G}_c$ (Eqn 1), of the transformation from local element coordinates $\{\zeta_\alpha\}$ to physical coordinates $\{x_i\}$. Shakib [65] considers the one dimensional limiting case of this multidimensional definition for the advection-diffusion equation and presents a comparison with the original SUPG technique,

$$[\mathbf{G}_c]_{ij} = \frac{\partial \zeta_\alpha}{\partial x_i} \frac{\partial \zeta_\alpha}{\partial x_j}. \tag{1}$$

Momentum	$\mathbf{F}_{m,i} = \int_\Omega \Phi \mathbf{R}_{m,i} d\Omega + \sum_e \int_{\Omega_e} \rho \tau_m^{\#} (\mathbf{u} \cdot \nabla \Phi) \mathbf{R}_{m,i} d\Omega$
Total Mass	$F_P = \int_\Omega \Phi R_P d\Omega + \sum_e \int_{\Omega_e} \rho \tau_m^{\#} \nabla \Phi \cdot \mathbf{R}_m d\Omega$
Thermal Energy	$F_T = \int_\Omega \Phi R_T d\Omega + \sum_e \int_{\Omega_e} \rho \hat{C}_p \tau_T^{\#} (\mathbf{u} \cdot \nabla \Phi) R_T d\Omega$
Mass Fraction for Species k	$F_{Y_k} = \int_\Omega \Phi R_{Y_k} d\Omega + \sum_e \int_{\Omega_e} \rho \tau_{Y_k}^{\#} (\mathbf{u} \cdot \nabla \Phi) R_{Y_k} d\Omega \quad , k = 1, 2, ..., N_s - 1$ $F_{Y_{N_s}} = R_{Y_{N_s}}$

Table 2
Stabilized finite element formulation of transport/reaction PDEs, where the residual equations R_i are presented in Table 1 and the stabilization parameters $\tau_i^{\#}$ are defined in Table 3.

Finally, it should be noted that this formulation, as presented, has some limitations. First this formulation is a simplification of proposed formulations for multiple advection-diffusion type equations that also couple the various equations in the definition of the least squares operators [38]. In addition there is no stabilization contribution for strong source terms and there is no nonlinear discontinuity capturing term present (see e.g.[38,71,65,13,7,21]). While we have experimented with these terms they are not used in the results presented in this overview. In the context of our overview, the current simplified formulation has the advantage of having a continuous stabilization operator. In addition it is possible to develop analytic Jacobians that are accurate for the limiting case of constant transport coefficients. These properties are useful for both for efficient implementation of the Newton-Krylov methods and effective preconditioning techniques, as well as for development of PDE constrained optimization techniques that require accurate Jacobians of the PDE problem [59].

2.2 Brief Overview of Discrete Systems of Equations

To give context to our later discussion of the solution methods and linear algebra, we present here a discussion of the structure of the equations that result from the FE discretization of the weak form of the transport/reaction equations. In this discussion the Newtonian stress tensor is expanded to in-

Momentum	$\tau_m^{\#} = (\rho^2 \mathbf{u}\mathbf{G}_c\mathbf{u} + 12\mu^2\|\mathbf{G}_c\|)^{-\frac{1}{2}}$
Thermal Energy	$\tau_T^{\#} = ((\rho C_p)^2 \mathbf{u}\mathbf{G}_c\mathbf{u} + 12\lambda^2\|\mathbf{G}_c\|)^{-\frac{1}{2}}$
Mass Fraction for Species k	$\tau_{Y_k}^{\#} = (\rho^2 \mathbf{u}\mathbf{G}_c\mathbf{u} + 12(\rho D_k)^2\|\mathbf{G}_c\|)^{-\frac{1}{2}}$

Table 3
Definition of stabilization parameters used in stabilized equations, which use the contravarient metric tensor $\mathbf{G}_c$ (Eqn 1) to define an element-level streamwise length scale. The superscript (#) denotes that this parameter has the units of time/density

clude the hydrodynamic pressure, P, and the viscous stress tensor term, Ψ . The resulting stabilized from of the total mass residual equation in expanded form is given by

$$F_P = \int_\Omega \Phi[\frac{\partial \rho}{\partial t} + \nabla \cdot (\rho \mathbf{u})]d\Omega + \sum_e \int_{\Omega_e} \rho\tau_m^{\#} \nabla\Phi \cdot [\rho\frac{\partial \mathbf{u}}{\partial t} + \rho\mathbf{u} \cdot \nabla\mathbf{u} + \nabla P - \nabla \cdot \mathbf{\Psi} - \rho\mathbf{g}]d\Omega. \quad (2)$$

This expansion includes the weak form of a Laplacian operator acting on pressure,

$$K = \sum_e \int_{\Omega_e} \rho\tau_m^{\#} \nabla\Phi \cdot \nabla P d\Omega, \quad (3)$$

which is produced by the stabilized formulation of the total mass conservation equation. Finite element (FE) discretization of the stabilized equations gives rise to a system of coupled, nonlinear, non-symmetric algebraic equations, the numerical solution of which can be very challenging. These equations are linearized using an inexact form of Newton's method as described in Section 3. A formal block matrix representation of these discrete linearized equations is given in Eqn 4 where the block diagonal contribution of the stabilization procedure has been highlighted by a specific ordering. In this representation, the vector, $\mathbf{v}'$, contains the Newton updates to all the nodal solution variables with the exception of the nodal pressures, $\mathbf{P}'$. The block matrix, $\mathbf{A}$, corresponds to the combined discrete convection, diffusion, and reaction operators for all the unknowns; the matrix, $\mathbf{B}$, corresponds to the discrete divergence operator with its transpose the gradient operator; the diagonal matrix, $\mathbf{R}$, results from the group FE expansion of the density and velocity; and the matrix, $\mathbf{K}$, corresponds to the discrete "pressure Laplacian" operator discussed above. The

vectors $\mathbf{F_v}$ and $\mathbf{F_P}$ contain the right hand side residuals for Newton's method.

The existence of the well conditioned nonzero matrix, $\mathbf{K}$, in the stabilized FE discretization of the equations allows the solution of the linear systems with a number of algebraic and domain decomposition type preconditioners [62,64]. This is in contrast to other formulations, such as Galerkin methods using mixed interpolation, that produce a zero block on the total mass continuity diagonal. The difficulty of producing robust and efficient preconditioners for the Galerkin formulation has motivated the use of many different types of solution methods. A number of these use two-level iteration schemes, penalty methods, pseudo-compressibility techniques or decoupled/segregated solvers (e.g. [74,75]). A detailed presentation of the characteristics of current solution methods is far beyond the scope of this brief overview. However the intent of our method of fully-coupling the transport PDEs in the nonlinear solver is to preserve the inherently strong coupling of the physics with the goal to produce a more robust solution methodology. Preservation of this strong coupling, however, places a significant burden on the linear solution procedure to solve the fully coupled algebraic systems,

$$\begin{bmatrix} \mathbf{A} & -\mathbf{B^T} \\ \mathbf{BR} & \mathbf{K} \end{bmatrix} \begin{bmatrix} \mathbf{v'} \\ \mathbf{P'} \end{bmatrix} = - \begin{bmatrix} \mathbf{F_v} \\ \mathbf{F_P} \end{bmatrix}. \tag{4}$$

Our current linear algebra solution procedure uses a specific ordering of the unknowns locally at each FE node with each degree of freedom order consecutively. A single coupled matrix problem, $\mathbf{Js} = -\mathbf{F}$, is solved at each Newton step with sophisticated algebraic domain decomposition and multilevel preconditioned Krylov methods as described in Section 3.3.

2.3 General Description of Evolution, Steady State, and Parameter Dependent Equations

For clarity of the discussions that follow, we present a compact notation for the functional dependence of the discretized residuals. The time dependence of the governing equations and the dependence on the physical parameters of the system are made explicit as follows. Formally we represent this nonlinear system of coupled equations in Table 2 as

$$\mathbf{F}(\dot{\mathbf{v}}, \mathbf{v}, \mathbf{p}) = \mathbf{0}, \tag{5}$$

where $\mathbf{F}$, is the vector of residuals for the stabilized FE equations, $\mathbf{v}$, is the vector of unknowns, and the list of parameters on which the system depends

is denoted by $\mathbf{p}$. In this paper we focus on steady solutions $(\mathbf{0}, \bar{\mathbf{v}}, \bar{\mathbf{p}})$ of Eqn 5 corresponding to solutions such that:

$$\mathbf{F}(0, \bar{\mathbf{v}}, \bar{\mathbf{p}}) = \mathbf{0}. \tag{6}$$

3 Obtaining Steady State Solutions: Parallel Newton-Krylov Solver

3.1 Problem partitioning

Chaco [26], a general graph partitioning tool, is used to partition the FE mesh into sub-domains and make sub-domain-to-processor assignments. Chaco constructs partitions and sub-domain mappings that have low communication volume, good load balance, few message start-ups and only small amounts of network congestion. Chaco supports a variety of new and established graph partitioning heuristics. For the results in this paper, multilevel methods with Kernighan-Lin improvement were used. For a detailed description of parallel FE data structures and a discussion of the strong link between partitioning quality and parallel efficiency, see [26,63].

3.2 Newton-Krylov Methods

As described above the stabilized FE formulation of the governing steady state transport equations form a strongly coupled nonlinear system of equations (Eqn 6). In our solution procedure we use an inexact Newton method with back-tracking [19,61] to solve this system. For purposes of continuity with the discussion that will follow on continuation and linear stability analysis, we present a Newton's method motivated by a Taylor series expansion about the current approximate steady solution $\mathbf{v}_{i-1}$, holding the parameters $p_{(i)}$ fixed,

$$\mathbf{F}(\mathbf{v}_k) = \mathbf{F}(\mathbf{v}_{k-1}) + \mathbf{F}_{\mathbf{v}}(\mathbf{v}_{k-1})[\mathbf{v}_k - \mathbf{v}_{k-1}] + O(\|\mathbf{v}_k - \mathbf{v}_{k-1}\|^2). \tag{7}$$

As is well known, Newton's method is recovered from Eqn 7 by setting $\mathbf{F}(\mathbf{v}_k) = \mathbf{0}$ and defining the correction vector $\Delta\mathbf{v}_k \equiv \mathbf{v}_k - \mathbf{v}_{k-1}$ and the Jacobian matrix $\mathbf{J}_{k-1} \equiv \mathbf{F}_{\mathbf{v}}(\mathbf{v}_{k-1})$. This produces the linear systems (Eqn 8) that are solved at each step of Newton's method,

$$\mathbf{J}_{k-1}\Delta\mathbf{v}_k = -\mathbf{F}(\mathbf{v}_{k-1}). \tag{8}$$

Since typical applications for 3D reacting flow solutions require the solution of linear systems with millions of equations and unknowns we rely on iterative

methods for the underlying linear system solvers. For this reason a Newton-Krylov method is usually implemented as an inexact Newton method [19,61]. That is, in approximately solving Eqn 8, one chooses a forcing term $\eta \in [0, 1)$ and then applies a Krylov method until an iterate satisfies the inexact Newton condition,

$$\|\mathbf{F}(\mathbf{v}_{k-1}) + \mathbf{J}_{k-1}\Delta\mathbf{v}_k\| \leq \eta\|\mathbf{F}(\mathbf{v}_{k-1})\|. \tag{9}$$

In this context an inexact Newton method uses nonlinear residual information to determine the accuracy η_k, to which the sequence of linear subproblems are solved. Intuitively one would assume that in the initial stages of the Newton iteration, when the current approximation is far from the true solution, there would be no benefit from solving the Newton equations to too high a degree of accuracy. The inexact Newton method uses an adaptive convergence criteria to reduce the amount of over-solving that occurs and thereby to produce a more computationally efficient nonlinear solution procedure. Back-tracking is a technique for improving the robustness of the nonlinear solver, and works by scaling the Newton correction vector as needed to ensure that the nonlinear residual has been reduced adequately before the step is accepted. Specific choices for η_k, in the inexact Newton scheme and a more thorough numerical evaluation of these methods can be found in [19,61,54].

3.3 Preconditioned Krylov Methods

The linear subproblems generated from the inexact Newton method are solved by preconditioned Krylov methods as implemented in the Aztec parallel iterative solver library [40]. The Krylov algorithms implemented in Aztec include techniques such as the restarted generalized minimal residual [GMRES(k)] and transpose-free quasi-minimal residual (TFQMR) techniques for non-symmetric systems. In this overview paper, we illustrate briefly the performance of restarted GMRES(k) and TFQMR with domain decomposition type preconditioners.These preconditioners are based on algebraic additive Schwarz domain decomposition (DD) methods with variable overlapping between sub-domains.

The overall performance of Krylov methods can be substantially improved when one uses preconditioning [56,2]. The basic idea is that instead of solving the system $Ax = b$, the system $AM^{-1}y = b$ is solved, where M^{-1} is an approximation to A^{-1} and is easily computed. Since only matrix-vector products are needed, it is not necessary to explicitly form AM^{-1} (only software to solve $Mv = y$ is needed).

3.3.1 Additive Schwarz Domain Decomposition Preconditioner

In this section, we define the variable overlap additive Schwarz DD preconditioners that are employed. A more detailed discussion of these methods can be found in [64] and a general reference is [68]. The i^{th} overlap additive two level Schwarz preconditioner is given by

$$\mathbf{M}^{-1} = \sum_{k=0}^{P} \mathbf{I}_k^i (\tilde{\mathbf{A}}_k^i)^{-1} (\mathbf{I}_k^i)^T, \tag{10}$$

where P is the number of processors (equivalent in our case to the number of sub-domains), $\mathbf{I}_0^i$ is an interpolation operator that maps solution vectors from an auxiliary coarser mesh – which covers the entire domain but with significantly fewer grid points – to the original fine mesh, and $(\tilde{\mathbf{A}}_0^i)^{-1}$ is the inverse of the discrete operator on the coarse mesh.

To complete the description we define the transfer operators and the discrete sub-domain operators, $\mathbf{I}_k^i$, for $k > 0$. $\mathbf{I}_k^i$ is a rectangular restriction matrix containing only zeros and ones. It injects solutions on the fine mesh to the vertices of the k^{th} sub-domain that has i levels of overlap. The vertices corresponding to the 0^{th} overlap of the k^{th} sub-domain correspond to the vertex set V_k^0 (which is produced by the Chaco decomposition) where

$$V_i^0 \bigcap V_j^0 = \emptyset \ \ \forall i \neq j \tag{11}$$

and $\bigcup_{i=1}^{P} V_i^0$ contains all vertices (FE nodes) in the graph (FE mesh). The vertices of the m^{th} overlap sub-domain are defined recursively by taking all vertices of distance one from V_k^{m-1} in the level set of the graph connectivity (i.e. the FE nodes adjacent to the current sub-domain). Finally, $\mathbf{A}_k^i = (\mathbf{I}_k^i)^T \mathbf{A} \mathbf{I}_k^i$ and $(\tilde{\mathbf{A}}_k^i)^{-1}$ is an approximation to $\mathbf{A}_k^i$'s inverse on the sub-domain by performing for example an incomplete LU factorization (ILU) or a Gauss-Seidel relaxation (smoothing) technique.

4 Mapping Complex Solution Spaces: Continuation and Bifurcation Analysis

The strength of the above robust solution strategy is not just in locating a single steady-state solution, but in the ability to perform engineering and scientific analysis. This involves the determination of the nonlinear behavior of the system with respect to key parameters. To this end, continuation techniques are used extensively [45,42,76] in conjunction with linear stability analysis techniques discussed in the following section.

Euler-Newton continuation is an effective method of continuing a regular path of points near a given regular point $(\mathbf{v}_0, p_0)$ [42]. To determine a neighboring point $(\mathbf{v}_1, p_1)$ corresponding to $p_1 = p_0 + \Delta p$, where Δp is a sufficiently small step in the physical parameter, the following predictor corrector strategy is used. First, the tangent vector $\mathbf{v}_p \equiv \frac{\partial \mathbf{v}}{\partial p_0}$ is computed from Eqn 6 by differentiating with respect to p,

$$\mathbf{F}_{\mathbf{v}}\mathbf{v}_{p_0} + \mathbf{F}_{p_0} = \mathbf{J}_0\mathbf{v}_{p_0} + \mathbf{F}_{p_0} = \mathbf{0}, \tag{12}$$

which can then be solved for $\mathbf{v}_{p_0}$ with one linear solve of the Jacobian. This would be the final result for a linear sensitivity analysis, but to track the nonlinear behavior one must continue. A prediction of $(\mathbf{v}_1, p_1)$ is then obtained by:

$$\mathbf{v}_1^{(1)} = \mathbf{v}_0 + \Delta p \mathbf{v}_{p_0}. \tag{13}$$

This is used as a starting point for Newton's method when solving

$$\mathbf{F}(\mathbf{v}_1, p_1) = \mathbf{0}. \tag{14}$$

This algorithm works well for regular solution paths, with the prediction (Eqn 13) enabling Newton's method to converge quickly. Since this technique is used successively to continue along the branch, the continuation approach can be effective when the solution of Eqn 8 can be obtained efficiently. Euler-Newton continuation clearly breaks down at simple limit points where the solution "turns back" on itself, where the choice of $p_1 = p_0 + \Delta p$ may constrain p_1 to be in a region where no solution $(\mathbf{v}_1, p_1)$ exists. However at these points the Implicit Function Theorem still holds and an alternate parametrization of the solution branch is possible.

In this case a pseudoarclength technique proceeds by selecting a step Δs in an "arc-length" like parameter [42], by which the solution $\mathbf{v}(s)$ and parameter $p(s)$ have been parameterized. The pseudoarclength technique is used to continue robustly through very steep changes of the solution with respect to p, and even past limit (or turning) points in the solution branch. The approach is to add an additional parameter, s, and an arc-length constraint $N(\mathbf{v}, p, s)$ to the equations governing the steady solutions of Eqn 6. That is, the system

$$\begin{aligned} \mathbf{F}(\mathbf{v}(s), p(s)) &= \mathbf{0}, \\ N(\mathbf{v}(s), p(s), s) &= 0, \end{aligned} \tag{15}$$

is used, where the pseudoarclength parameter, s, parameterizes the solution branch. At simple turning points, where s continues to increase monotonically

even though p does not, the procedure can be proven to continue past the singular point for sufficiently small Δs [42].

The coupled system in Eqn 15 is solved by application of Newton's method. Newton's method generates a block matrix system of equations as given in Eqn 16,

$$\begin{bmatrix} \mathbf{J} & \mathbf{F}_p \\ \frac{\partial \mathbf{v}^T}{\partial s} & \frac{\partial p}{\partial s} \end{bmatrix} \begin{bmatrix} \Delta \mathbf{v} \\ \Delta p \end{bmatrix} = - \begin{bmatrix} \mathbf{F}(\mathbf{v}, p) \\ N(\mathbf{v}, p, s + \Delta s) \end{bmatrix}. \tag{16}$$

The solution of Eqn. 16 can either be again carried out by Newton-Krylov methods on this augmented problem, or can be decomposed into 2 individual solves with $\mathbf{J}$. These algorithms are implemented in the general purpose Library of Continuation Algorithms (LOCA) software package as described in [57].

Locating bifurcations is important for engineering analysis because they often correspond to the onset of undesirable behavior, such as the onset of oscillations, the breaking of symmetry, or ignition/extinction behavior. At bifurcation points the Implicit Function Theorem does not hold and multiple solution branches intersect. Once a bifurcation point is located and the eigenvalue and corresponding eigenvector associated with the bifurcation point have been calculated, the LOCA library can locate a bifurcation point and track it as a function of a second parameter. Routines have been implemented in LOCA to track turning point, pitchfork, and Hopf bifurcations. The tracking algorithms continue in a parameter, $p_{(1)}$ (called the continuation parameter), and calculate the value of a second parameter, $p_{(2)}$ (called the bifurcation parameter) such that the steady state solution remains on the bifurcation point. In this way, two-parameter plots (a.k.a. bifurcation sets), which map the boundaries between qualitatively different solution regimes, can be directly generated. The curves generated by these algorithms often delineate good and bad designs.

As an example, the tracking algorithm for the pitchfork bifurcation requires the solution of the following system of equations:

$$\mathbf{F}(\mathbf{x}, p_{(1)}, p_{(2)}) + \varepsilon \mathbf{w} = 0, \tag{17}$$

$$\mathbf{x} \cdot \mathbf{w} = 0, \tag{18}$$

$$\mathbf{J}\mathbf{n} = 0, \tag{19}$$

$$\mathbf{w} \cdot \mathbf{n} - 1 = 0. \tag{20}$$

These $2m + 2$ equations are solved for the solution vector $\mathbf{x}$, the null vector $\mathbf{n}$,

the bifurcation parameter $p_{(2)}$, and a slack parameter ε representing the asymmetry in the system, which typically goes to zero. The initial guess for the null vector, $\mathbf{n}^{(0)}$, is computed with the eigensolver at a point near the bifurcation. The vector $\mathbf{w}$ is antisymmetric with respect to the symmetry being broken, which in practice is chosen as $\mathbf{n}^{(0)}$. The system of equations requires that there is a steady state solution, that the solution is on the symmetric branch, and that the Jacobian has a null vector which is non-trivial. Solving this system of $2m+2$ equations and unknowns repeatedly while performing continuation in $p_{(1)}$ will produce the desired locus of symmetry breaking bifurcations in 2-parameter space.

In [57] algorithms have been developed which solve this nonlinear system using just solves of the Jacobian matrix $\mathbf{J}$. Similar algorithms have been developed for the two generic instabilities seen in 1-parameter systems, turning points (folds) and Hopf bifurcations. With this approach, sophisticated design and analysis capabilities are available to codes that can robustly solve the linear systems with the same Jacobian matrix as in the Newton step in Eqn 8.

5 Determining Stability of a Solution Branch: Linear Stability Analysis

Using the continuation procedures of Section 4, the multiplicity of steady solutions available to the system Eqn 6 can be explored. In addition to the existence of such multiple solutions, it is of equal interest to determine the stability or long time behavior of such solutions when subjected to various disturbances. The response of the physical system to infinitesimal disturbances is determined by a linearizion of the nonlinear system in Eqn 5. The background and theory of linear stability methods is well developed and widely used [50,11]. The appropriate linearizion of Eqn 5 is formally produced with a multi-variate Taylor series expansion about the steady solution at $(\mathbf{0}, \bar{\mathbf{v}})$:

$$\mathbf{F}(\dot{\mathbf{v}}, \mathbf{v}) = \mathbf{F}(\mathbf{0}, \bar{\mathbf{v}}) + \mathbf{F}_{\dot{\mathbf{v}}}[\dot{\mathbf{v}} - \mathbf{0}] + \mathbf{F}_{\mathbf{v}}[\mathbf{v} - \bar{\mathbf{v}}] + \mathit{Higher\ Order\ Terms}. \quad (21)$$

Neglecting higher order terms, using Eqn 5 and Eqn 6 and defining the generalized mass matrix, $\mathbf{B} \equiv -\mathbf{F}_{\dot{\mathbf{v}}}$, and $\mathbf{v}' \equiv \mathbf{v} - \bar{\mathbf{v}}$ as the deviation from the steady solution, we obtain the evolution equations for infinitesimal disturbances,

$$\mathbf{B}\dot{\mathbf{v}}' = \mathbf{J}\mathbf{v}'. \quad (22)$$

The disturbance quantity $\mathbf{v}'$, is given by

$$\mathbf{v}' = \sum_{i=1}^{N} \epsilon_i e^{\sigma_i t} \mathbf{q}_i, \tag{23}$$

where $\mathbf{q}_i$ is a generalized eigenvector (or normal mode) with eigenvalue σ_i. Eqn 22 then can be written as

$$\sigma_i \mathbf{B}\mathbf{q}_i = \mathbf{J}\mathbf{q}_i. \tag{24}$$

Equation 24 is the generalized eigenvalue problem governing stability of the linearization of the evolution equations at the steady solution $(\mathbf{0}, \bar{\mathbf{v}})$.

A steady state is linearly stable when all the amplitude coefficients of Eqn 23 decay in time, i.e. when all the eigenvalues have negative real part. A steady state is linearly unstable when one or more of the eigenvalues has positive real part. In practice, the eigenvalues of algebraically largest real part in the spectrum of the linearization are monitored as one of the parameters, say p^1, is varied. As an eigenvalue crosses the imaginary axis, the steady state of the nonlinear system becomes unstable to infinitesimal disturbances. At a simple bifurcation point the crossing eigenvalue has zero imaginary part, from such a point a steady solution branches. The continuation of this branch is accomplished by the methods of Section 4. The linear stability of the steady solution branch is determined as described above. The situation for nonzero imaginary part is however, different. According to the Hopf bifurcation theorem, a branch of time periodic solutions branches from a steady solution if a complex conjugate pair of eigenvalues of the Jacobian matrix crosses the imaginary axis.

As described computational linear stability analysis calculations rely on the solution of the generalized eigenvalue problem (24). To compute the rightmost eigenvalues, a generalized Cayley spectral transform [51,46] is used to reformulate the generalized eigenvalue problem into an ordinary eigenvalue problem [47,49]:

$$\mathbf{T}_s \mathbf{w} = (\mathbf{J} - \eta_c \mathbf{B})^{-1} (\mathbf{J} - \mu_c \mathbf{B}) \mathbf{w} = \nu \mathbf{w}, \quad \nu = \frac{\sigma - \mu_c}{\sigma - \eta_c}. \tag{25}$$

Here η_c is the pole and μ_c is the zero of the Cayley transform, and ν is the transformed eigenvalue. Selection of the Cayley parameters, η_c and μ_c, is critical to map the correct set of eigenvalues (those with the largest real part) to the eigenvalues of the largest magnitude in the Cayley transform. The importance of this transform is that it maps the infinite negative eigenvalues

generated by the continuity equation (resulting in a singular $\mathbf{B}$ from the incompressibility assumption) to a value of one. This makes the spectral condition number of the system smaller than other typical transformations such as the shift-invert spectral transformation, and therefore generally easier to apply the inverse operator with iterative linear solvers [49]. The resulting ordinary eigenvalue problem is solved using an implicitly restarted Arnoldi's method in P_ARPACK [48,49,58,8].

6 Results and Discussion

6.1 An Illustration of Linear Solver Performance and Scaling

In this section we present some representative results for a family of Schwarz domain decomposition (DD) methods [9,10,68]. These schemes partition the original domain into subdomains and approximately solve the discrete problems corresponding to the individual subdomains in parallel. Among Schwarz schemes, there are a number of choices which can greatly affect the overall solution time and robustness. These choices include the subdomain size, the amount of overlap between subdomains, and the partitioning metric which can alter the shape and aspect ratio of subdomains (see e.g.[22,68,27,28,64]). The choices also include the selection of subdomain solver such as an incomplete LU factorization (ILU) (with further options for dropping nonzeros in the factorizations and ordering equations within a subdomain [2]), and the introduction of a coarse grid solve [23,72]. In our overview we briefly comment on the benefit from employing a coarse grid solve in a two-level DD method.

To compare the performance of the one-level and two-level Schwarz domain decomposition preconditioners, a scalability study is presented for a standard two-dimensional benchmark thermal convection flow problem [15] along with a three-dimensional version of this problem. This example problem models a thermal convection (or buoyancy-driven) flow in a differentially heated square box in the presence of gravity (Figure 1). The temperature on the heated wall and other parameters are chosen so that the Rayleigh number is 1000. The three-dimensional problem adds two no-slip insulated walls in the third dimension to form a cube of unit dimensions. These simple geometries facilitate scalability studies as different mesh sizes can be easily generated. The results were obtained on the ASCI-Red Tflop computer at Sandia National Laboratories. Each node of this machine contains 256 MB RAM and 333 MHz Pentium II processors. For the one-level DD preconditioner, an incomplete factorization (ILU) was used. For the two-level preconditioner, the fine mesh smoother was two sweeps of Gauss-Seidel, and the coarse mesh direct solver was SuperLU [16]. The fine mesh is generated by uniform refinement of the

coarse mesh. A fine-to-coarse grid ratio of roughly 64 in two dimensions and 512 in three dimensions was used.

As an illustration of parallel efficiency we consider the scaling of the one-level DD ILU preconditioner as presented in Table 4 and 5. The Krylov method is a transpose free quasi-minimum residual (TFQMR) technique. The results indicate that the two main kernels of the steady state solver, the matrix fill (or Jacobian creation) and the preconditioned iterative solver (TFQMR with DD ILU) scale very well. The matrix creation phase which is based on local element integration is nearly perfectly scalable, and the linear solver shows a very high degree of scaled efficiency as well up to 2048 processors in the case of the 3D simulations.

Fig. 1. 2D and 3D benchmark thermal convection problem, $Ra = 10^3$. Figures show color contours of temperature on the plane and in 3D positive (blue) and negative velocity (yellow) isosurfaces and streamlines.

Num. proc	mesh	Num. unknowns	avg iter /Newt step	total time (sec)	avg time /matrix fill (sec)	scaled eff.	avg time /linear iter (sec)	scaled eff.
32	33^2	4356	76	1.5	0.0246	—	0.00279	—
128	65^2	16900	193	5.5	0.0244	1.01	0.00297	0.93
512	129^2	66564	696	12	0.0242	1.02	0.00327	0.84

Table 4
Efficiency of one-level DD-ILU preconditioner for 2D thermal convection problem with TFQMR

The effects of the parallel and algorithmic scaling of the one- and two-level preconditioners for the two- and three-dimensional thermal convection problems are presented in Tables 6 and 7. Figure 2 summarizes the results for the average iteration count per Newton step as a function of problem size. As the number of unknowns, N (as well as the number of processors, P, in this scaled study), is increased the number of iterations to convergence for the one-level schemes increases significantly: roughly $N^{2/3}$ in two dimensions and $N^{1/2}$ in three dimensions. Note that an optimal convergence property is obtained for the two-level preconditioner. The parallel version of SuperLU was used to factor the coarse matrix in groups of approximately $P^{1/2}$ processors. Since the

Num. proc	mesh	Num. unknowns	avg iter /Newt step	total time (sec)	avg time /matrix fill (sec)	scaled eff.	avg time /linear iter (sec)	scaled eff.
32	33^3	179,685	105	279	6.297	—	0.324	—
256	65^3	1,373,125	324	680	6.276	0.99	0.345	0.94
2048	129^3	10,733,445	860	1749	6.271	1.01	0.364	0.89

Table 5
Efficiency of one-level DD-ILU preconditioner for 3D thermal convection problem with TFQMR

proc	fine grid size	fine grid unknowns	1-level ILU		2-level GS-SuperLU			
			avg its/ Newt step	time (sec)	coarse grid	coarse unks size	avg its/ Newt step	time (sec)
1	33^2	4356	41	23	5^2	100	32	18
4	65^2	16,900	98	62	9^2	324	33	26
16	129^2	66,564	251	275	17^2	1156	34	30
64	257^2	264,196	603	1399	33^2	4356	31	46
256	513^2	1,052,676	1478	8085	65^2	16900	30	107

Table 6
Comparison of 1-level and 2-level scheme for 2D thermal convection problem

proc	fine grid size	fine grid unknowns	1-level ILU		2-level GS2-SuperLU			
			avg its/ Newt step	time (sec)	coarse grid size	coarse unks	avg its/ Newt step	time (sec)
4	17^3	24,565	40	123	3^3	135	36	101
32	33^3	179,685	112	282	5^3	625	44	107
256	65^3	1,373,125	296	863	9^3	3645	47	179
2048	129^3	10,733,445	650	2915	17^3	24565	47	546

Table 7
Comparison of 1-level and 2-level scheme for 3D thermal convection problem

fine grid smoother is highly parallel and the fine grid work per processor is roughly constant, the cost of the direct solve (SuperLU) on the increasingly larger coarse grid causes an increase in the CPU time for the larger problems.

While this loss of CPU time scaling is disturbing it must be pointed out that the two-level methods are still significantly faster (a factor of about 80 in 2D and 6 in 3D) than the corresponding one-level methods. To mitigate this growth of CPU time for the coarse grid solve either approximate coarse grid methods can be used (e.g. [64]) or more levels could be employed. In addition, the ability of the two-level methods to provide a sufficient linear residual decrease in the sub-problems for Newton's method allows for a more robust iterative nonlinear solver. We have seen many examples for which a direct-to-steady-state nonlinear calculation will fail, for very ill-conditioned large linear systems, because the iterative linear solver converges too slowly (to be practical) or does not converge at all. Further details of the parallel performance and algorithmic scalability of these methods can be found in [63,64].

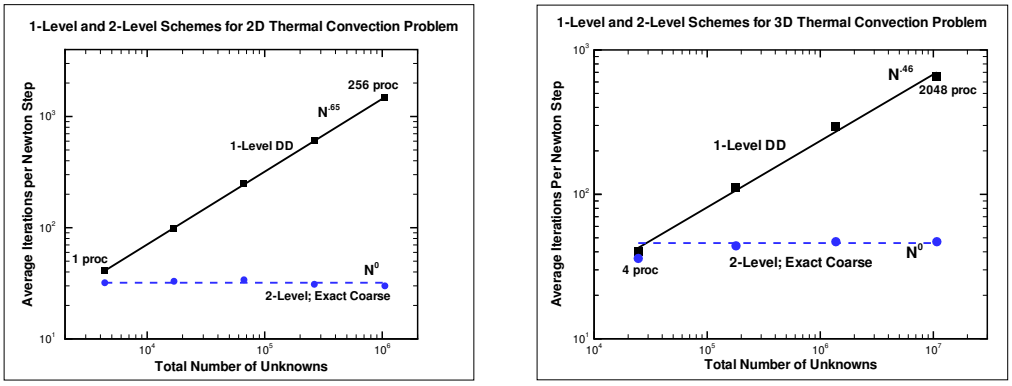

Fig. 2. Scaling of iteration count for one- and two-level preconditioners for 2D and 3D thermal convection problem

In Figure 3 we attempt to interpret the effect of the coarse grid correction for the two-level scheme physically. This figure presents the x-component of velocity on a slice plane through the center of the three-dimensional thermal convection problem for a sequence of meshes from 2^3 hexahedral elements to 64^3 hexahedral elements. The figures in the left, center and right columns represent the coarse and fine mesh in the first, second and third rows respectively of Table 7. For the two figures in the left column, the 2^3 element coarse mesh does not really resolve any basic (long wavelength or coarse) features of the flow field, so it is no surprise that the two-level preconditioner reduced the number of iterations by only 10%. For the two figures in the middle column, the 4^3 element coarse mesh starts to resolve the basic features of the flow field, which is why the two-level preconditioner reduced the number of iterations by 60%. For the two figures in the right column, the 8^3 element coarse mesh does a significantly better job at resolving the features of the flow field than the 4^3 element coarse mesh, hence the significant reduction in iterations by a factor of 6 over the one-level preconditioner. In the case of the 2048 processor problem (10+ million unknowns) this reduction is larger than a factor of 13.

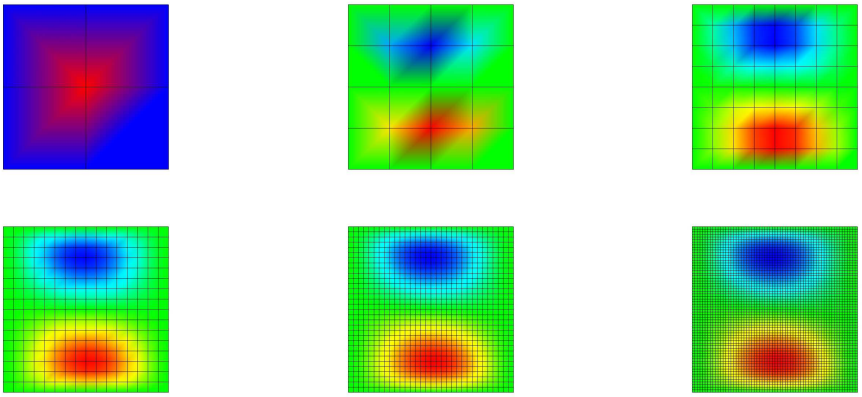

Fig. 3. x-component of velocity on a slice plane through the center of the 3D thermal convection problem for a sequence of meshes

6.2 An Illustration of Mapping Complex Nonlinear Solution Spaces: Bifurcation and Stability Analysis

6.2.1 Impinging Jet Reactor Example [53]

Symmetric counter-flowing jets have been used extensively for studies involving endothermic decomposition kinetics, diffusion flame kinetics, polymer processing, micro/nano-particle synthesis, blood flow, and other applications. A comprehensive discussion of impinging jet reactor applications can be found in [69].

Understanding the structure and stability of these flows is critical in interpreting experimental results in such reactors. In a collaboration with Prof. T.J. Mountziaris of SUNY at Buffalo, the algorithmic tools outlined in Sections 4 and 5 were used to perform a stability and bifurcation analysis of a laminar isothermal impinging jet reactor [53]. This section includes selected results from that work to illustrate the power of combining robust solution algorithms with sophisticated analysis capabilities.

The initial model consists of two isothermal jets with equal mass flow rates. Figure 4 shows a side view schematic of the impinging jet reactor used in this study. Two dimensionless parameters dictate the behavior of the impinging jets: the Reynolds number based on the jet separation distance and the aspect ratio, defined as jet width compared to the jet separation distance. The procedure for analyzing a two-parameter system involves first performing continuation runs on one parameter while monitoring the stability with the linear stability analysis capability. When a bifurcation is detected, the bifurcation tracking algorithms are then used to track the critical parameter value corresponding to the bifurcation point as a function of the second parameter. Using a relatively small number of simulations, a design map can be constructed where qualitatively different flows are delineated by the bifurcation

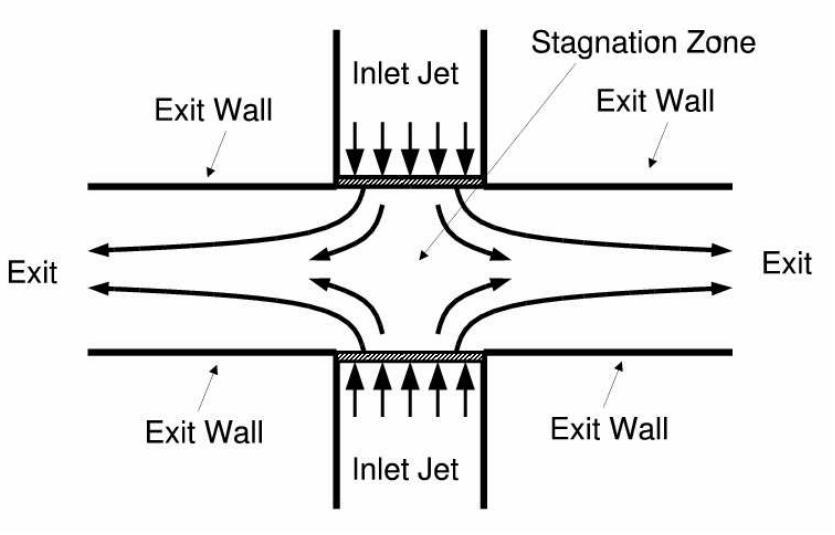

Fig. 4. Side view of an impinging jet reactor. Both cartesian and cylindrical models are considered in this work.

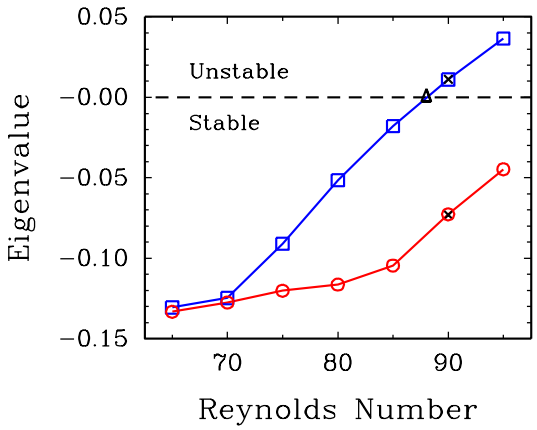

Fig. 5. Plot of the two leading eigenvalues as a function of jet Reynolds number, for the 1.7M unknown mesh. Bifurcation is as the leading eigenvalue crosses from negative to positive. The triangle data point is a bifurcation point located by the bifurcation tracking algorithm and the two cross data points are the eigenvalues calculated on a mesh of 10M unknowns.

points as a function of the parameters of interest.

At low Reynolds numbers, a single symmetric steady state was found to exist. The flow forms a well defined stagnation zone with large convectively driven recirculations on each side of the inlet jets. For certain geometric configurations, increasing the Reynolds number led to a loss of stability via a pitchfork bifurcation. This was identified by the rightmost eigenvalue being real-valued and passing through 0 as the Reynolds number is increased. Figure 5 shows the leading two eigenvalues, which happen to be real, plotted as a function of Reynolds number (increasing in fixed steps of 5.0) for an aspect ratio of 1.0. The exact bifurcation point was pinpointed with the pitchfork tracking algo-

rithm. The results were performed on a 3D mesh of 1.7 million unknowns on 24 (3.06 GHz Pentium Xeon) processors. The linear solver was a single-level scheme, with an ILUT preconditioned GMRES. A typical continuation step took 3 Newton iterations and 8 minutes to solve with 96% of the time in the linear solver and 4% in the matrix fill algorithm. The linear stability analysis took another 30 minutes per step. Superimposed results on a 10 million unknown mesh at a Reynolds number of 90 showed that the eigenvalues are well converged with mesh. The eigenvalue calculations for this mesh required about 45 minutes on 200 Processors. The lack of algorithmic scalability (increasing linear iteration count) on this problem underlies the motivation for multilevel preconditioning schemes.

The (super-critical) pitchfork bifurcation leads to multiplicity in the steady-state solution with three steady states existing for each Reynolds number above the critical bifurcation point value. An example of the three steady states are shown in Figure 6 for a fixed aspect ratio and a Reynolds number of about 10% above the critical value. The two rightmost eigenvalues are listed

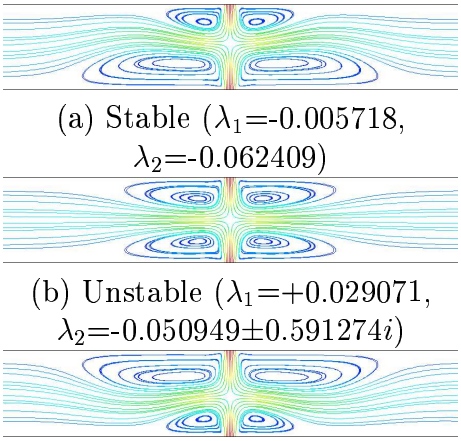

(a) Stable (λ_1=-0.005718, λ_2=-0.062409)

(b) Unstable (λ_1=+0.029071, λ_2=-0.050949$\pm$0.591274i)

(c) Stable (λ_1=-0.005718, λ_2=-0.062409)

Fig. 6. Streamlines for rectangular impinging jets at a Reynolds number of 27.5 and aspect ratio of 0.125. Three steady state solutions are present. (a) stable asymmetric upper branch, (b) unstable symmetric branch, and (c) stable asymmetric lower branch. The two rightmost eigenvalues are listed for each solution.

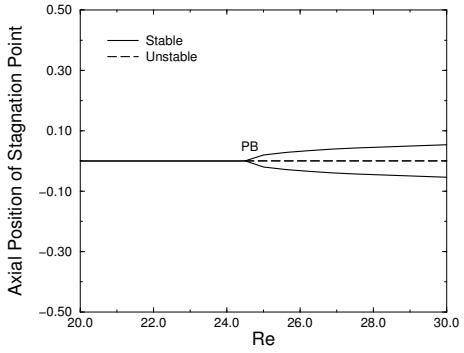

Fig. 7. Bifurcation diagram at aspect ratio of 0.125. The position of the stagnation point on the reactor centerline (z-axis) is plotted against the Reynolds number. The upper and lower inlet jets are positioned at -0.5 and 0.5, respectively. The label PB is the pitchfork bifurcation point. Solid lines are stable steady states, dashed lines are unstable steady states.

for each steady state. The real parts of the rightmost eigenvalues of the two asymmetric solutions, (a) and (c), are negative indicating that they are stable steady states. The symmetric solution, (b), has a positive real part indicating that it is unstable to small disturbances.

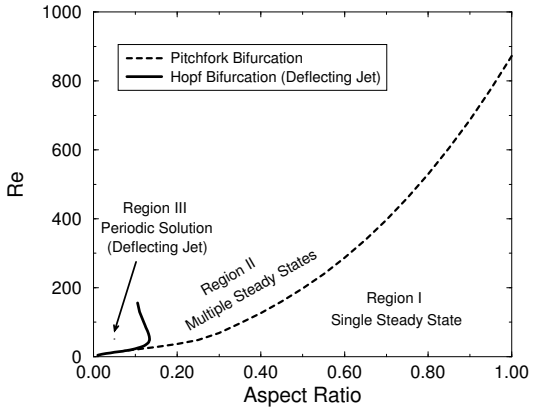

Fig. 8. Two-parameter bifurcation set showing different flow regimes, delineated by pitchfork and Hopf bifurcations.

Using the continuation routines, we were able to track each of the steady state solutions and construct a bifurcation diagram. The symmetry breaking of the pitchfork bifurcation shifts the stagnation point in the vertical direction. The extent of the shift of the stagnation point is shown in the bifurcation diagram in Figure 7. As one can imagine, the qualitatively different flow solution on different sides of the bifurcation point can significantly alter the performance or reliability of the system.

For certain conditions in a rectangular coplanar geometry, we found in reference [53] that the single stable symmetric jet solution can lose symmetry to an oscillatory Hopf bifurcation. This is detected by a complex conjugate pair of eigenvalues moving across the imaginary axis as the Reynolds number is increased. This flow mode is described as a "deflecting jet oscillation" and is the behavior that Denshchikov et al. observed and traced out experimentally [17,18].

For the geometric configurations studied above, the pitchfork and Hopf bifurcations each serve to separate qualitatively different flow regimes. By using the bifurcation tracking algorithms these results can be generalized to a range of aspect ratios. The results of bifurcation tracking runs are the curves in Figure 8, the two-parameter bifurcation set for this problem. It provides a cohesive picture for impinging jet physics with rectangular inlets. The symmetric stagnation flow Region I is seen to exist for all aspect ratios, while Region II with the asymmetric steady states and Region III with the deflecting jets oscillations only exist for certain aspect ratios. The leading eigenvalues along the bifurcation curve are monitored to confirm that these modes of instability are not overtaken by other modes, which does occur at the higher-codimension bifurcation where the Hopf and pitchfork curves intersect.

The stability analysis routines have also been employed on a non-isothermal chemically reacting model that simulates non-premixed laminar diffusion flames

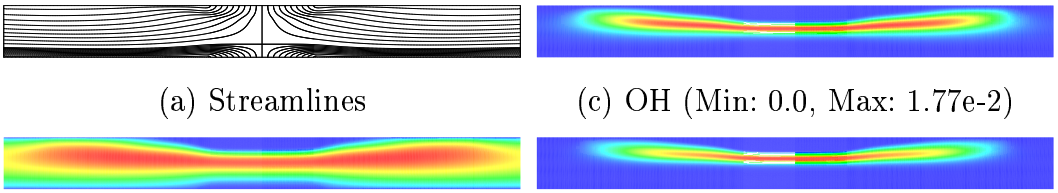

(a) Streamlines (c) OH (Min: 0.0, Max: 1.77e-2)

(b) Temp. (Min: 300^{o}K, Max: 2727^{o}K) (d) O (Min: 0.0, Max: 5.9e-3)

Fig. 9. Hydrogen/Oxygen diffusion flame in a non-premixed counterflow jet reactor. Species contours for OH and O are mass fraction values.

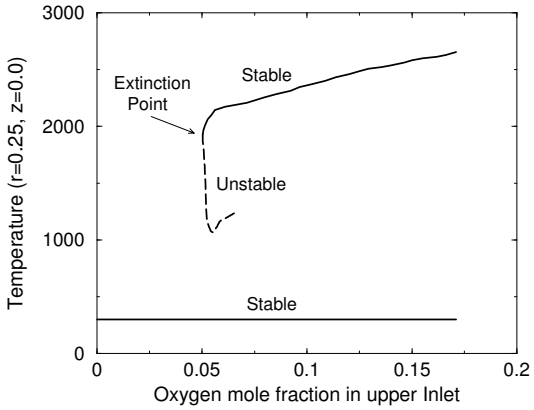

Fig. 10. Plot of the turning point (extinction point) of the temperature (at the point z=0.0, r=0.25cm) as a function of O_2 concentration of the upper oxidizer stream in the counterflow jet reactor. The lower non-ignited branch is also depicted.

in counter-flowing jets. In this 2D axisymmetric model, the upper inlet jet is the oxidizer stream, consisting of oxygen diluted in argon (O_2, and Ar) and the lower inlet jet is the fuel stream, consisting of hydrogen diluted argon (H_2, and Ar). A reaction mechanism developed by Conaire et al. [14] employs 10 chemical species (we ignore the inert He species in the original mechanism) and 19 gas phase reactions. Figure 9 shows a plot of the fluid streamlines, along with contour plots of the temperature, hydroxyl and oxygen atom mass fraction intermediates which indicate the location of the flame. The fluid enters the reactor at 300^{o}K and has a maximum value of 2727^oK inside the flame. This increase in temperature across the flame of almost a factor of 10 accounts for the acceleration of the oxidizer and the deflection of the streamlines downward.

Using the arc-length continuation routine, the extinction point of the diffusion flame was located. Figure 10 is a bifurcation diagram plotting the temperature at a fixed point in the reactor (vertically centered and shifted radially by 0.25cm) as a function of the oxygen mole fraction in the upper jet. Three solution branches were identified, an upper stable branch where the reactor is in the ignited state, a middle unstable branch still in the ignited state, and a lower stable branch where the flow is non-ignited. The upper stable and

middle unstable branches of solutions are connected through a turning point. This turning point is the extinction point. Reducing the oxygen mole fraction in the upper jet below this value moves the model into an oxygen-starved regime where only one solution exists, the non-ignited branch.

The computational performance for solving the 14 coupled transport/reaction PDEs by the standard one-level DD method is presented in Table 8. Timings and the scaled efficiency are shown for a single continuation step along the path to ignition of the gas streams. These results again indicate high parallel scalability for the local element integrations and the chemical kinetics evaluations that make up the matrix fill (Jacobian creation). The linear solver efficiencies are also very good as well.

Num Procs	level of refine.	Num. unknowns	avg time /matrix fill (sec)	scaled eff.	avg time /linear iter (sec)	scaled eff.
16	0	199,374	13.94	—	0.5219	—
64	1	790,734	13.86	1.01	0.5286	0.99
256	2	3,149,454	14.06	0.99	0.5363	0.97
1024	3	12,570,894	13.99	1.00	0.5369	0.97

Table 8
Scaled Efficiency of hydrogen/oxygen diffusion flame simulation with DD-ILU TFQMR.

6.3 An illustration of Solution of Complex Transport/Reaction Systems

6.3.1 Partial Catalytic Oxidation of Ethane

This example presents representative performance and simulation results for catalytic partial oxidation of ethane to ethylene. Ethylene is one of the most widely used chemicals, with the US producing 25 billion kilograms annually. Catalytic partial oxidation (CPO), a very rapid and autothermal process, would provide a huge capital cost saving over the traditional steam cracking process [3]. The goal of this study was to look for processing conditions where the selectivity and conversion of CPO are competitive of those for steam cracking. The standard production process of steam cracking produces ∼85% selectivity with a conversion of ∼60% of ethane. Using Platinum-coated foam monoliths, Huff and Schmidt [29,30] found that a selectivity of ∼65% to ethylene could be achieved at a conversion of ∼70%. Computational investigation of alternative operating conditions, particularly at high pressures, is preferable to experimental studies due to safety concerns.

The reactor under consideration is a straight tube monolith (a bundle of straight tube channels with the same inlet feed). In this study, we look at tubes that are 4 cm long: a 1 cm non-catalytic inlet section (radiation shield), followed by a 1 cm section coated with Platinum, followed in turn by a 2 cm non-catalytic outlet section. The inside diameter of the tube is 0.6 mm. The computational domain is one tube, modeled with 2D cylindrical coordinates, and includes heat conduction in the wall so that the natural autothermal operation can be simulated.

The reaction mechanism was supplied by Dave Zerkle of Los Alamos National Labs and comes from his collaboration with Mark Allendorf of Sandia National Labs. The mechanism consists of 22 gas-phase species with 77 reversible gas-phase reactions, and 17 surface species undergoing 35 surface reactions. The CHEMKIN and SURFACE CHEMKIN packages are used to provide rigorous treatment of the multicomponent transport properties and species production/consumption rates. All physical properties in the gas-phase are treated as functions of the local temperature and mole fractions and a global operating pressure.

Modeling this reactor presents numerical challenges. The ethane/air mixture passes through the reactor at over 1 meter per second and the temperature in the reactor changes by a factor of four within millimeters, both producing very large gradients. The reaction terms (both surface and gas phase) and the dependence of the physical properties on the local state cause significant non-linearities. Scaling is also an issue, since some of the intermediate species never reach mole fractions above 10^{-5} yet are significant to the reactive process.

Solutions to the governing set of 26 coupled PDEs were obtained using the Finite Element code MPSalsa developed by Sandia National Laboratories. The calculations presented here were on a fine mesh of 2727 nodes corresponding to 70902 unknowns. The Aztec linear solver package is used, and an ILUT preconditioner with overlap and fill-in and the GMRES iterative solver were selected. A pseudo-time step time integration process was needed once to get to a first steady state solution. Once a steady state solution was achieved, all other solutions were calculated by steady state parameter continuation.

Here, we will present results from parameter studies on two key parameters, the inlet feed ratio of carbon to oxygen ratio (C_2H_6/O_2) and the reactor pressure. The effects on performance on four other other parameters have been studied in a technical report [52].

Selected solution variables at steady state are shown in Figure 11. The flow is from left to right, the axis of symmetry is at the top, and the bottom red strip is the solid wall, where only the heat conduction equation is solved.

The velocity plot shows the acceleration due to the increase in temperature

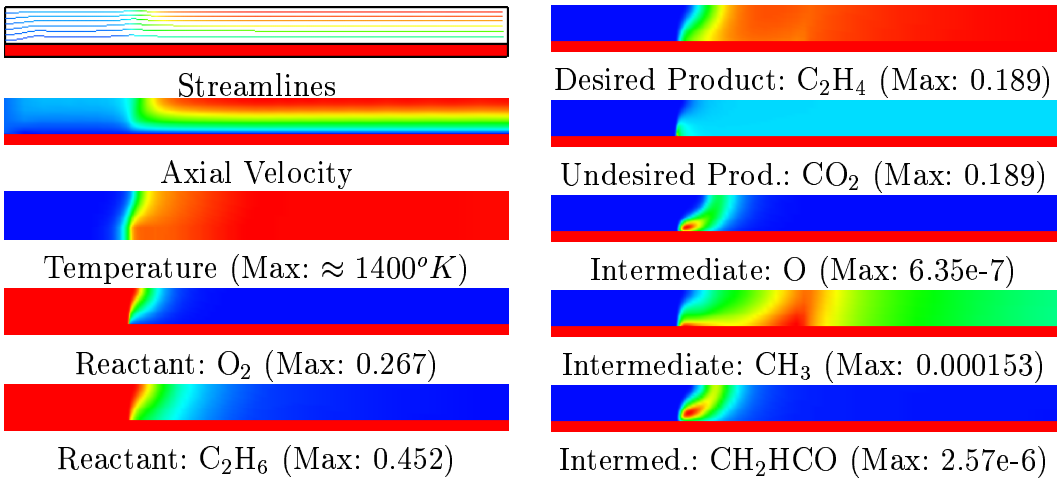

Fig. 11. Contour plots of 10 of the 22 solution variables at steady state. For visualization purposes, the radial (vertical) dimension on the the plots has been stretched a factor of 10. Only the temperature is solved for in the wall, which appears red in the other plots. The maximum mass fraction achieved in the reactor appears for the species unknowns.

and the net decrease in the mixture molecular weight of the products, while the temperature plot shows the large gradient near the beginning of the catalyst. The plots of the reaction intermediates show how different elementary reactions take place in different zones of the reactor. The most relevant plot is of CO_2 mass fraction, the undesirable product of total oxidation, which is produced exclusively in a very sharp peak right at the beginning of the catalyst. Design alternatives that restrict the availability of oxygen at this point may significantly increase the selectivity to C_2H_4.

A parametric study was undertaken, using the algorithms presented in Section 4, to identify the optimal operating conditions for reactor based upon the currently proposed catalysis and chemistry mechanisms. The first parameter of interest is the carbon to oxygen (C/O) ratio to verify the computations against the experimental observations from Veser and Schmidt [73]. Figure 12 depicts the conversion and selectivities for an increasing C/O ratio from 1.8 to 2.4. Note that there is an extreme drop in conversion (down to 65%) with a marginal increase in ethylene selectivity. The C/O parameter is critical in maximizing the efficiency of this reactor system.

Figure 13 shows the conversion and selectivity plots for a continuation study of the reactor pressure parameter. The interesting behavior is that the selectivity goes through a maximum around 4.5 atm. Overall, the performance is a weaker function of pressure then the C/O ratio.

The computational analysis of the catalytic partial oxidation reactor has helped to elucidate the dependence of key parameters on the performance of the reactor [52]. Robustness and efficiency of the steady-state solution pro-

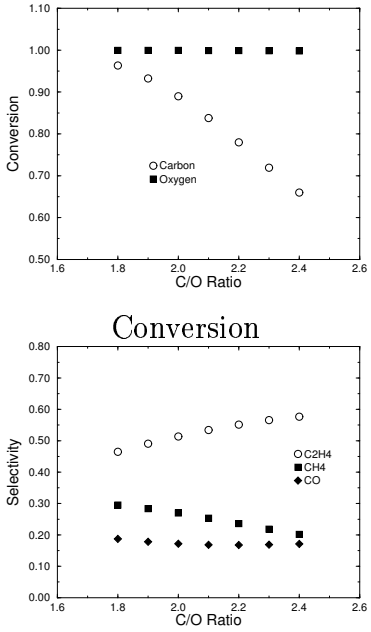

Conversion

Carbon Selectivity

Fig. 12. Conversion and selectivity of Carbon atoms as a function of inlet carbon to oxygen ratio, as calculated from a continuation run.

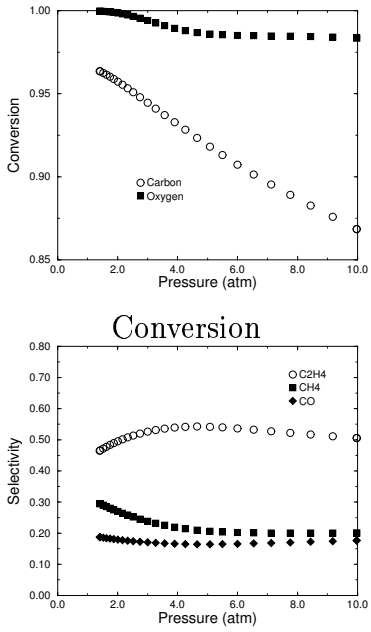

Conversion

Carbon Selectivity

Fig. 13. Conversion and selectivity of Carbon atoms as a function of reactor pressure, as calculated from a continuation run.

cedure is crucial for performing parametric studies and, as will be shown in the following example, even performing optimization of complex reacting flow systems.

6.3.2 Chemical Vapor Deposition of Poly-Silicon

This second example presents representative CPU time performance for the one- and two-level preconditioners applied to chemical vapor deposition (CVD) of poly-Silicon. Figure 14 shows the geometry and steady-state calculation for a three-dimensional reacting flow simulation for the deposition of poly-silicon in a horizontal rotating disk reactor. A mixture of trichlorosilane ($SiCl_3H$), HCl, and H_2 enters from the four inlets on the left, flows over a forward facing step, and over an inset rotating disk. The disk is heated to 1398K, which initiates chemical reactions to deposit silicon on the wafer. The chemical reaction model of Kommu, Wilson, and Khomami [44] was used, along with a similar reactor design. This calculation is a full reacting flow problem with unknowns for the three velocity components, hydrodynamic pressure, temperature, and the above three chemical species, and nonlinear surface reaction boundary conditions.

Figure 14 shows the schematic picture of the reactor, along with representative streamlines through the reactor and deposition rate contours along the

rotating susceptor. Optimization calculations have been performed on this re-

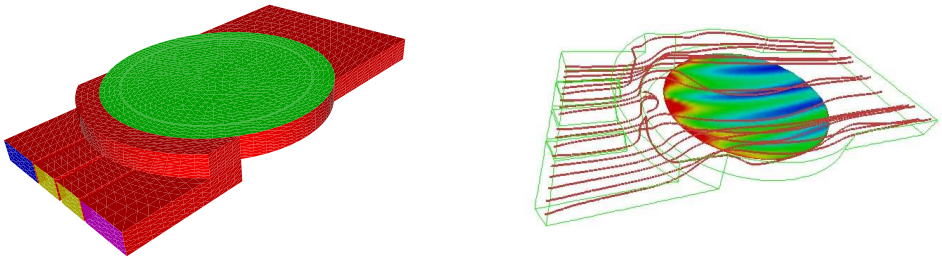

Fig. 14. 3D horizontal epitaxial-silicon CVD reactor, including 4 split inlets and step; right image shows flow streamlines and a contour plot of poly-silicon deposition rate on the susceptor surface (flow is from left to right, and red represents high deposition rates)

actor to minimize the non-uniformity of the deposition rate over the disk, after the averaging effect of rotation is factored in. Instantaneous deposition rate profiles are shown in Figure 15 and show the remnants of the split inlets,

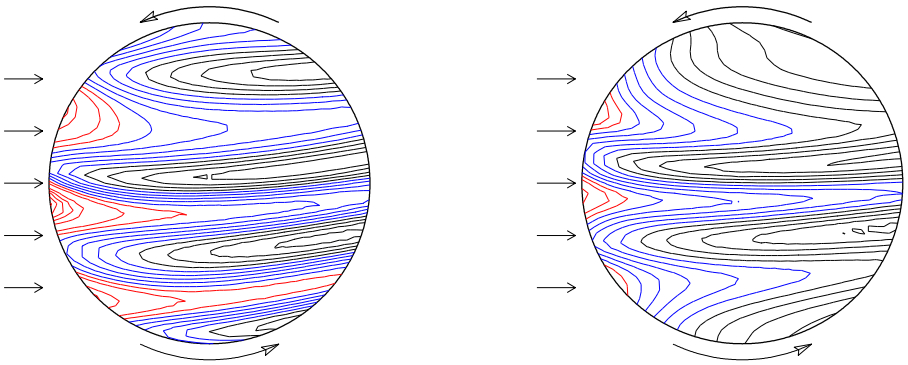

Fig. 15. Contour plot of instantaneous deposition of poly-silicon in 3D horizontal CVD reactor with split inlets; left image shows initial non-optimized flow/transport/reaction conditions and right image shows optimized flow/transport/reaction conditions over 4 parameters.

though this effect is decreased in the optimized solution. The radial deposition profiles for initial parameter values and the results of 2 optimization calculations are shown in Figure 16. The standard deviation of these profiles is the objective function minimized by the optimization algorithm over key operating parameters. Clearly, robust and scalable solution algorithms are required for these complex and large systems, since they form the inner kernels of the optimization runs.

Representative performance results are presented from a simple continuation step where the reactor thermodynamic pressure was increased from 0.6 atmospheres to the operating pressure of 0.85 atmospheres. Preliminary scaling results are presented in Table 9 for a sequence of hexahedral meshes. The first four rows of the table present a scaled study where each successive row is a uniform refinement of the mesh in the previous row. The number of unknowns

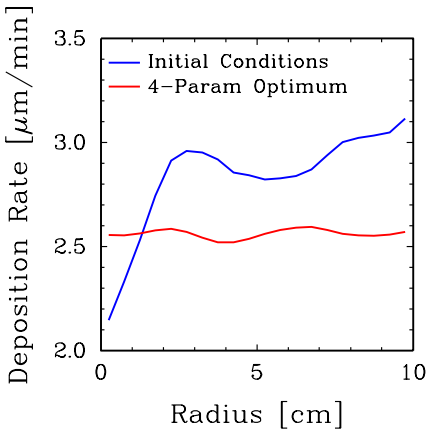

Fig. 16. Radial profiles of angularly-averaged deposition (due to disk rotation) of Poly-Si in 3D horizontal CVD reactor. Optimization runs minimize the non-uniformity of the deposition rates. Plot shows initial non-optimized, and the 4-parameter optimized conditions (optimization variables: inlet velocity profiles, wafer rotation rate, the purge flow velocity and inlet reactant concentration profiles)

is varied from 87,000 to 38 million. The solution method was non-restarted GMRES with ILU as the preconditioner with the convergence criterion for the linear solve was chosen to be 3×10^{-4}. These calculations were performed on the Sandia Cplant machine, which consists of 500 MHz Dec Alpha processors with Myrinet interconnect (performance of the 500 MHz Dec Alphas were very roughly comparable to 1 GHz Pentium 3 CPUs).

The final row of the table presents results with the 2-level preconditioner with ILU smoother on the fine mesh and ILU approximate solve on the coarse mesh of 87,000 unknowns. These preliminary results show that the two-level preconditioner for the largest problem was about 25% faster than the one-level preconditioner. While this reduction in time is not as impressive as for the fluid flow calculations presented earlier, there is still a reasonable benefit to using the two-level method. Note that this is a particularly difficult problem to solve and a direct-to-steady-state fully-coupled solver is being used.

Calculations with the two-level preconditioner were not performed for any meshes other than the finest. This is because a geometric two-level preconditioner requires that a coarse mesh be generated. The 87,000 unknown mesh is the coarsest mesh generated for this problem, and the cost of this coarse mesh solve with a fine mesh that is only about eight times larger would be very significant compared to the fine mesh. In practice, the fine mesh needs to be substantially larger than the coarse mesh to obtain benefits from the two-level preconditioner. For the 37.8M unknown problem, the coarse mesh of 87,000 unknowns was large enough to preclude using a direct solver, which is why the approximate ILU solve was used. Use of a multilevel preconditioner (more than two levels) would alleviate these constraints, as very coarse meshes

proc	level of ref	unknowns	Preconditioner	avg iter/ Newt step [steps]	total time (sec)
2	0	87,000	DD-ILU	47 [6]	742
16	1	636,000	DD-ILU	88 [8]	1531
128	2	4,846,000	DD-ILU	200 [9]	3703
1000	3	37,807,000	DD-ILU	482 [10]	10973
1000	3	37,807,000	2 level (ILU/ILU)	230 [10]	8174

Table 9
One- and two-level preconditioner for 3D horizontal epitaxial-silicon CVD reactor

could be obtained. Work on a multilevel algebraic preconditioner for reactive flow problems is currently being pursued.

7 Conclusions

This paper has presented a brief discussion of a stabilized FE formulation for low Mach number flow, thermal energy transfer and mass species transport with non equilibrium chemical reactions. An overview of the stabilized FE formulation was described along with the numerical solution methods required for detailed analysis of complex steady state transport/reaction systems. These solution algorithms include robust nonlinear and linear solution schemes, parameter continuation methods, and linear stability analysis techniques. Using these techniques and the computational power of massively parallel computing we have demonstrated the ability to carry out detailed and efficient engineering analysis of complex nonlinear transport/reaction systems with the goal of producing an understanding of the physical mechanisms which stabilize and destabilize such systems.

8 Acknowledgments

The authors wish to thank Scott Hutchinson, Karen Devine, and Harry Moffat for their collaborative effort in developing the MPSalsa code. We would also like to thank Tayfun Tezduyar for early discussions on stabilized FE methods and Farzin Shakib for many helpful discussions and notes on the development of stabilized FE methods for incompressible flow with thermal transport. Finally we would like to thank Prof. T.J. Mountziaris of the University of Buffalo

for collaborating on the impinging jet reactor problem.

References

[1] T. Barth, P. Bochev, M. Gunzburger, and J. Shadid. A taxonomy of consistently stabilized finite element methods for the Stokes problem. *SIAM J. Sci. Comp.*, 25(5):1585–1607, 2004.

[2] M. Benzi. Preconditioning techniques for large linear systems: A survey. *J. Comp. Physics*, 182:418–477, 2002.

[3] A. S. Bodke, D. A. Olschki, L. D. Schmidt, and E. Ranzi. High selectivities to ethylene by partial oxidation of ethane. *Science*, 258:712–715, 1999.

[4] F. Brezzi. On existence, uniqueness and approximation of saddle-point problems arising from Lagrange multipliers. *RAIRO Model. Math. Anal. Numer.*, 21:129–151, 1974.

[5] A. N. Brooks and T.J.R. Hughes. Streamline upwind/petrov-galerkin formulations for convection dominated flows with particular emphasis on the incompressible navier-stokes equations. *Comp. Meth. Appl. Mech. and Eng.*, 32:199–259, 1982.

[6] P. N. Brown and Y. Saad. Convergence theory of nonlinear Newton–Krylov algorithms. *SIAM J. Optimization*, 4:297–330, 1994.

[7] E. Burman and A. Ern. Nonlinear diffusion and discrete maximum principle for stabilized galerkin approximations of the convection-diffusion-reaction equation. *Comp. Meth. Appl. Mech. and Eng.*, 191:3833–3855, 2002.

[8] E.A. Burroughs, R.B. Lehoucq, L.A. Romero, and A.G. Salinger. Linear stability of flow in a differentially heated cavity via large-scale eigenvalue calculations. *Int J Numer Meth Heat Fluid Flow*, 14(6):803–822, 2004.

[9] X.-C. Cai. An additive schwarz algorithm for nonselfadjoint elliptic equations. In T.F. Chan, Roland G., Jacques P., and O. B. Widlund, editors, *Third International Symposium on Domain Decomposition Methods for Partial Differential Equations*, pages 232–244, 1989.

[10] X.-C. Cai, W. D. Gropp, and D. E. Keyes. Convergence rate estimate for a domain decomposition method. *Numer. Math.*, 61(2):153–169, 1992.

[11] S. Chandrasekhar. *Hydrodynamic and Hydromagnetic Stability.* Oxford Univ. Press, 1961.

[12] K. A. Cliffe, A. Spence, and S. J. Tavener. The numerical analysis of bifurcation problems with applciation to fluid mechanics. *Acta Numerica*, pages 39–131, 2000.

[13] R. Codina. Comparison of some finite element methods for solving the diffusion-convection-reaction equations. *Comp. Meth. Appl. Mech. and Eng.*, 156:185–210, 1998.

[14] M.O. Conaire, H.J. Curran, J.M. Simmie, W.J. Pitz, and C.K. Westbrook. A comprehensive modeling study of Hydrogen oxidation. *Int. J. Chem. Kinet.*, 2004. in press.

[15] G.D. Davis and I.P. Jones. Natural convection in a square cavity: A comparison exercise. *Int. J. Numer. Methods Fluids*, 3(3):227–248, 1983.

[16] J. W. Demmel, J. R. Gilbert, and X. S. Li. Superlu users' guide. Technical Report LBNL-44289, Lawrence Berkeley National Laboratory, September 1999.

[17] V. A. Denshchikov, V. N. Kontratev, and A. N. Romashev. Interaction between two opposed jets. *Fluid Dynamics*, 6:924–926, 1978.

[18] V. A. Denshchikov, V. N. Kontratev, A. N. Romashev, and V. M. Chubarov. Auto-oscillations of planar colliding jets. *Fluid Dynamics*, 3:460–462, 1983.

[19] S. C. Eisenstat and H. F. Walker. Globally convergent inexact Newton methods. *SIAM J. Optimization*, 4:393–422, 1994.

[20] S.C. Eisenstat and H.F. Walker. Globally convergent inexact Newton methods. *SIAM J. Optimization*, 4:393–422, 1994.

[21] A. Ern and J.L. Guermond. *Theory and practice of finite elements.* Springer-Verlag, New York, 2004.

[22] C. Farhat, N. Maman, and G.W. Brown. Mesh partitioning for implicit computations via iterative domain decomposition - impact and optimization of the subdomain aspect ratio. *Int. J. Num. Meth. Eng.*, 38(6):989–1000, 1995.

[23] W.D. Gropp, D.K. Kaushik, D.E. Keyes, and B.F. Smith. High-performance parallel implicit cfd. *Parallel Computing*, 27:337–362, 2001.

[24] J. Guckenheimer and P. Holmes. *Nonlinear oscillations, dynamical systems, and bifurcations of vector fields.* Springer-Verlag, 1983.

[25] M. Gunzburger. *Finite Element Methods for Viscous Incompressible Flows.* Academic Press, Boston, 1989.

[26] B. Hendrickson and R. Leland. An improved spectral graph partitioning algorithm for mapping parallel computations. Technical Report SAND 92-1460, Sandia National Laboratories, Albuquerque, NM, 1992.

[27] B. A. Hendrickson. Load balancing fictions, falsehoods and fallacies. *Applied Mathematical Modelling*, 25(2):99–108, 2000.

[28] B. A. Hendrickson and T. G. Kolda. Graph partitioning models for parallel computing. *Parallel Computing*, 26(12):1519–1534, 2000.

[29] M. Huff and L. D. Schmidt. Ethylene formation by oxidative dehydrogenation of ethane over monoliths at very short contact times. *Journal of Physical Chemistry*, 97:11815, 1993.

[30] M. Huff and L. D. Schmidt. Elementary step model of ethane oxidative dehydrogenation on pt-coated monoliths. *AIChE Journal*, 42(12):3484–3497, 1996.

[31] T. J. R. Hughes, L. P. Franca, and G. M. Hulbert. A new finite element formulation for computational fluid dynamics: VII. the Galerkin/Least-Squares method for advective-diffusive equations. *Computer Methods in Applied Mechanics and Engineering*, 73:173–189, 1989.

[32] T. J. R. Hughes and T.E. Tezduyar. Finite element methods for first-order hyperbolic systems with particular emphasis on the compressible euler equations. *Comp. Meth. Appl. Mech. Engrg.*, 45:217–284, 1984.

[33] T.J.R. Hughes. Multiscale phenomena: Green's functions, the dirichlet-to-neumann formulation, subgrid scale models, bubbles and the origins of stabilized methods. *Comp. Meth. Appl. Mech. and Eng.*, 127:387–401, 1995.

[34] T.J.R. Hughes and A. Brooks. A theoretical framework for Petrov-Galerkin methods with discontinuous weighting functions: Application to the streamline-upwind procedure. In R.H. Gallagher et al, editor, *Finite Elements in Fluids*, volume 4, pages 47–65. J. Willey & Sons, 1982.

[35] T.J.R. Hughes and A. N. Brooks. A multidimensional upwind scheme with no cross-wind diffusion. *in T.J.R. Hughes (ed.), Finite Element Methods for Convection Dominated Flows, AMD Vol. 34, ASME, New York 19-35*, 1979.

[36] T.J.R. Hughes and L. Franca. A new finite element formulation for computational fluid dynamics: VII. The Stokes problem with various well-posed boundary conditions: symmetric formulations that converge for all velocity pressure spaces. *Comput. Meth. Appl. Mech. Engrg.*,, 65:85–96, 1987.

[37] T.J.R. Hughes, L. Franca, and M. Balestra. A new finite element formulation for computational fluid dynamics: V. Circumventing the Babuska-Brezzi condition: A stable Petrov-Galerkin formulation of the Stokes problem accommodating equal-order interpolations. *Comput. Meth. Appl. Mech. Engrg.*, 59:85–99, 1986.

[38] T.J.R. Hughes and M. Mallet. A new finite element formulation for computational fluid dynamics: III. The generalized streamline operator for multidimentional advective-diffusive systems. *Comput. Meth. Appl. Mech. Engrg.*, 58:305–328, 1986.

[39] T.J.R. Hughes, M. Mallet, and A. Mizukami. A new finite element formulation for computational fluid dynamics: II. Beyond SUPG. *Comput. Meth. Appl. Mech. Engrg.*, 54:341–355, 1986.

[40] S.A. Hutchinson, L. Prevost, J.N. Shadid, C. Tong, and R.S. Tuminaro. Aztec user's guide version 2.0. Technical Report SAND 99-8801J, Sandia National Laboratories, 1999.

[41] R.J. Kee, F.M. Rupley, and J.A. Miller. Chemkin-II: A Fortran chemical kinetics package for the analysis of gas-phase chemical kinetics. Technical Report SAND 89-8009, Sandia National Laboratories, 1989.

[42] H.B. Keller. Numerical solution of bifurcation and nonlinear eigenvalue problems. In Rabinowitz, editor, *Application of Bifurcation Theory*, pages 45–52, Dekker, NY, 1977.

[43] D. A. Knoll and D. E. Keyes. Jacobian-free Newton–Krylov methods: a survey of approaches and applications. *J. Comput. Phys.*, to appear.

[44] S. Kommu, G. M. Wilson, and B. Khomami. A theoretical/experimental study of silicon epitaxy in horizontal single-wafer chemical vapor deposition reactors. *J. Elec. Chem. Soc.*, 147(4):1538–1550, 2000.

[45] M. Kubicek and M. Marek. *Computational Methods in Bifurcation Theory and Dissipative Structures.* Computational Physics. Springer-Verlag, New York, NY, 1983.

[46] R. B. Lehoucq and K. Meergbergen. Using generalized Cayley transformations within an inexact rational Krylov sequence method. *SIAM J. Matrix Analysis and Applications*, 20(1):131–148, 1998.

[47] R. B. Lehoucq and A. G. Salinger. Massively parallel linear stability analysis with P_ARPACK for 3D fluid flow modeled with MPSalsa. *Lecture Notes in Computer Science*, 1541:286–295, 1998.

[48] R. B. Lehoucq, D. C. Sorensen, and C. Yang. *ARPACK USERS GUIDE: Solution of Large Scale Eigenvalue Problems with Implicitly Restarted Arnoldi Methods.* SIAM, Phildelphia, PA, 1998.

[49] R.B. Lehoucq and A.G. Salinger. Large-scale eigenvalue calculations for stability analysis of steady flows on massively parallel computers. *Int. J. Num. Meth. Fluids*, 36:309–327, 2001.

[50] C.C. Lin. *The Theory of Hydrodynamic Stability.* Cambridge Univ. Press, 1955.

[51] K. Meerbergen, A. Spence, and D. Roose. Shift-invert and Cayley transforms for the detection of rightmost eigenvalues of nonsymmetric matricies. *BIT*, 34:409–423, 1994.

[52] R. P. Pawlowski and A. G. Salinger. Numerical simulation of the partial catalytic oxidation of ethane to ethylene in short contact time reactors. Technical Report Sand2001-1338, Sandia National Laboratories, Albuquerque NM, 87185, 1996.

[53] R. P. Pawlowski, A. G. Salinger, Shadid J. N., and Mountziaris T. J. Stability analysis of laminar isothermal impinging jet flows. *submitted to Journal of Fluid Mech.*, 2004.

[54] R. P. Pawlowski, J. N. Shadid, J. P. Simonis, and H. F. Walker. Globalization techniques for Newton-Krylov methods and applciations to the fully-coupled solution of the Navier-Stokes equations. *submitted to SIAM Review*, 2004.

[55] Y. Saad. Krylov subspace methods for solving large unsymmetric linear systems. *Mathematics of Computation*, 37(155):105–126, 1981.

[56] Y. Saad. *Iterative Methods for Sparse Linear Systems.* PWS Publishing Company, Boston, 1996.

[57] A. G. Salinger, E. A. Burroughs, R. P. Pawlowski, E. T. Phipps, and L. A. Romero. Bifurcation tracking algorithms and software for large scale applications. *Int. J. Bifurcation Chaos*, 2004. in press.

[58] A. G. Salinger, R. B. Lehoucq, R. P. Pawlowski, and J. N. Shadid. Computational bifrucation and stability studies of the 8:1 thermal cavity problem. *Int. J. Num. Meth. Fluids*, 40:1059–1073, 2002.

[59] A.G. Salinger, R.P. Pawlowski, J.N. Shadid, and B.G. van Bloemen Waanders. Computational analysis and optimization of a chemical vapor deposition reactor with large-scale computing. *Ind. Eng. Chem. Res.*, 43:4612–4623, 2004.

[60] J. N. Shadid, A. G. Salinger, R. C. Schmidt, S. A. Hutchinson, G. L. Hennigan, K. D. Devine, and H. K. Moffat. MPSalsa: A finite element computer program for reacting flow problems part 1: Theoretical development. Technical Report Sand98-2864, Sandia National Laboratories, Albuquerque NM, 87185, January. 1999.

[61] J. N. Shadid, R. S. Tuminaro, and H. F. Walker. An inexact Newton method for fully-coupled solution of the Navier–Stokes equations with heat and mass transport. *J. Comput. Phys.*, 137:155–185, 1997.

[62] J.N. Shadid. A fully-coupled Newton-Krylov solution method for parallel unstructured finite element fluid flow, heat and mass transfer simulations. *Int. J. CFD*, 12:199–211, 1999.

[63] J.N. Shadid, S.A. Hutchinson, G.L. Hennigan, H.K. Moffet, K.D. Devine, and A.G. Salinger. Efficient parallel computation of unstructured finite element reacting flow solutions. *Parallel Computing*, 23:1307–1325, 1997.

[64] J.N. Shadid, R.S. Tuminaro, K.D. Devine, G.L. Hennigan, and P.T. Lin. Performance of fully-coupled domain decomposition preconditioners for finite element transport/reaction simulations. *to appear in JCP*, 2005.

[65] F. Shakib. *Finite element analysis of the compressible Euler and Navier-Stokes equations.* PhD thesis, Division of Applied Mathematics, Stanford University, 1989.

[66] F. Shakib. Personal communication, 1997.

[67] F. Shakib, T.J.R. Hughes, and Z. Johan. A new finite element formulation for computational fluid dynamics: X. the compressible Euler and Navier-Stokes equations. *Comp. Meth. Appl. Mech. Engrg.*, 89:141–219, 1991.

[68] B. Smith, P. Bjorstad, and W. Gropp. *Domain Decomposition: Parallel Multilevel Methods for Elliptic Partial Differential Equations.* Cambridge University Press, 1996.

[69] Abraham Tamir. *Impinging Stream Reactors, Fundamentals and Applications.* Number 7 in Transposrt Processes in Engineering. Elsevier Science, 1994.

[70] T.E. Tezduyar. Stabilibized finite element formulations for incompressible flow calculations. *Advances in App. Mech.*, 28:1–44, 1992.

[71] T.E. Tezduyar and Y. J. Park. Discontinuity capturing finite element formulations for nonlinear convection-diffusion-reaction problems. *Comp. Meth. Appl. Mech. Engrg.*, 59:307–325, 1986.

[72] R.S. Tuminaro, C.H. Tong, J.N. Shadid, K.D.Devine, and D.M. Day. On a multilevel preconditioning module for unstructured mesh Krylov solvers: two-level schwarz. *Comm. Num. Method. Eng.*, 18:383–389, 2002.

[73] G. Veser and L. D. Schmidt. Ignition and extinction in the catalytic oxidation of hydrocarbons over platinum. *AIChE Journal*, 42(4):1077–1087, 1996.

[74] C. Vincent and R. Boyer. A preconditioned conjugate gradient uzawa-type method for the solution of the Stokes problem bu mixed Q1-P0 stabilized finite elements. *Int. J. Numer. Meth. Fluids*, 14:289–298, 1992.

[75] S.O. Wille. A preconditioned alternating inner-outer solution method for the mixed finite element formulation of the Navier-Stokes equations. *Int. J. Numer. Meth. Fluids*, 18:1135–1151, 1994.

[76] K.H. Winters, K.A. Cliffe, and C.P. Jackson. The prediction of instabilities using bifurcation theory. In R.W. Lewis, E. Hinton, P. Beltess, and B.A. Schrefler, editors, *Numerical Methods in Transient and Coupled Systems*, pages 179–198, 1987.